Homework Assignment 04
Technical Topic 01: Content Delivery Networks

January 20, 2010

F. Leighton and D. Lewin, *"Global Hosting System"*

This project had been issued as US Patent in 2000. Being a patent, it suggests a novel solution for distributed content delivery. By the time this patent was published, common way for content distribution was full replication of the content to different servers situated in various geographical locations. This was not effective, though, because of low cost efficiency and lack of techniques for centralized coordination of these machines. The system proposed suggests decentralized hosting solution, which allows users to obtain Internet content more efficiently and at the same time providers could manage the distribution of their content. The company which currently operates under this patent, and which founders are the authors of the present work, is Akamai.

Upon users' request the basic HTML document is generated by the content delivery server. The embeded objects in this HTML are then provided by the hosting servers (ghosts). To account for this, the authors implement threefold URL generation for each peace of content which is to appear in the HTML delivered to the end user. I find this approach for content allocation very successful. An URL consists of hosting server plus serial number of the content plus a fingerprint, which not only allows precise positioning but also delivery of the most up-to-date content. It also ensures more contol of the provider over the content that is distributed. A weakness I could emphasize here is that the time for generation of the basic HTML will increase as the number of embeded objects increase. It might be the case that different objects reside on different ghosts. Does this increase the time for which the HTML is bult and delivered? Overall, how cost effective is this approach from time prospective and do the authors make a tradeoff between users satisfaction and cheaper content replication. In this way of thoughts, a comparison of the time cost for basic HTML generation in conventional and in the proposed system would have been appropriate.

The system suggested here performs replication based on the number of requests for certain piece of content. Thus only parts of the whole content are being distributed and stored on servers with different geographical locations, and this decision is based on the level of user demand. One of the arguments that authors had against the common server mirroring is that in such setting

the providers have to contract with local hosting providers. In their new solution they also rely on content replication (although it's a smaller part of the entire content). But supposingly, the hosting facility contracting is still an issue here. In case it is not, what is their solution for storing whatever needs to be replicated?

When few ghosts are capable to provide the desired content, the framework should have such functionality, so to be able to choose which is the closest server to the user. To achieve this goal the system relies on DNS records providing information for the physical location of the ghost. Although this is a good idea for positioning, I believe it could suffer one problem - in order for proper DNS resolving there need to be only the IP address and the name of the hosts inscribed. Being optional, the physical address record is often ommited. This sadly might prevent the location algorithm to operate properly.

The location determinition of a ghost is not the only positioning problem in the proposed system. An other such is the location of the requesting user. After all, the optimality of content delivery relies on this information first. Relying on IP address lookup to determine physical location is not always the most appropriate solution. The handicaps here stems from the fact that some providers, being allocated IP address pools, attempt to further re-allocate slices to their customers, which might not always reside in the same geographical locations. Having said so, I think that there could be errors in positioning of the users, which also would result in un-optimal delivery.

An idea that I find quite successful is the improvement of the performance of long downloads by dynamical change the server from which the client downloads, based on network conditions. This would account for prompt delivery. I'd like to know, though, to what extent this promptness is compensated by the time needed for transition from one server to another.

As organized the paper is easy to read. It provides very good explanation of the figures applied - the numbering of the different components within a picture makes it way much easier for the reader (and coherent for the writer) to reference them in the process of explanation.

Being a patent, the present work is different as compared to the papers read so far. It actually does not provide evaluation of the proposed solution but only draws the concepts of what is to be done, for acheiving the ultimate goal - a framework for distributed content delivery. If the idea for publishing a patent is straightforward explanation of the novelty of the solution and what are the currently persisting problems that it solves, I'd definitely rather see it released as a conference paper, providing performance evaluation and comparison with the relevant to that time techniques for content delivery.